

Read Online Data Warehouse Design Solutions

Your all-in-one resource for using Pentaho with MySQL forBusiness Intelligence and Data Warehousing Open-source Pentaho provides business intelligence (BI) and datawarehousing solutions at a fraction of the cost of proprietary solutions. Now you can take advantage of Pentaho for your business needs with this practical guide written by two major participants in the Pentaho community. The book covers all components of the Pentaho BI Suite. You'll learn to install, use, and maintain Pentaho—and find plenty of background discussion that will bring you thoroughly up to speed on BI and Pentaho concepts. Of all available open source BI products, Pentaho offers the most comprehensive toolset and is the fastest growing open source product suite Explains how to build and load a data warehouse with Pentaho Kettle for data integration/ETL, manually create JFree (Pentaho Reporting Services) reports using direct SQL queries, and create Mondrian (Pentaho analysis services) cubes and attach them to a Pivot cube browser Review deploying reports, cubes and metadata to the Pentaho platform in order to distribute BI solutions to end-users Shows how to set up scheduling, subscription and automatic distribution The companion Web site provides complete source code examples, sample data, and links to related resources.

Over the past 5 years, the concept of big data has matured, data science has grown exponentially, and data architecture has become a standard part of organizational decision-making. Throughout all this change, the basic principles that shape the architecture of data have remained the same. There remains a need for people to take a look at the "bigger picture" and to understand where their data fit into the grand scheme of things. Data Architecture: A Primer for the Data Scientist, Second Edition addresses the larger architectural picture of how big data fits within the existing information infrastructure or data warehousing systems. This is an essential topic not only for data scientists, analysts, and managers but also for researchers and engineers who increasingly need to deal with large and complex sets of data. Until data are gathered and can be placed into an existing framework or architecture, they cannot be used to their full potential. Drawing upon years of practical experience and using numerous examples and case studies from across various industries, the authors seek to explain this larger picture into which big data fits, giving data scientists the necessary context for how pieces of the puzzle should fit together. New case studies include expanded coverage of textual management and analytics New chapters on visualization and big data Discussion of new visualizations of the end-state architecture

[Computer Networks and Information Technologies](#)

[The Big Book](#)

[Solutions for Star Schema Performance](#)

[Data Warehouses and OLAP](#)

[Building a Data Warehouse](#)

[The Data Warehouse Toolkit](#)

[Pentaho Solutions](#)

[Data Integration Best Practice Techniques and Technologies](#)

[Data Warehouse Systems](#)

[Impossible Data Warehouse Situations](#)

[Data Warehouse Designs](#)

[Design and Implementation](#)

DW 2.0: The Architecture for the Next Generation of Data Warehousing is the first book on the new generation of data warehouse architecture, DW 2.0, by the father of the data warehouse. The book describes the future of data warehousing that is technologically possible today, at both an architectural level and technology level. The perspective of the book is from the top down: looking at the overall architecture and then delving into the issues underlying the components. This allows people who are building or using a data warehouse to see what lies ahead and determine what new technology to buy, how to plan extensions to the data warehouse, what can be salvaged from the current system, and how to justify the expense at the most practical level. The book gives experienced data warehouse professionals everything they need in order to implement the new generation DW 2.0. It is designed for professionals in the IT organization, including data architects, DBAs, systems design and development professionals, as well as data warehouse and knowledge management professionals. * First book on the new generation of data warehouse architecture, DW 2.0. * Written by the "father of the data warehouse". Bill Inmon, a columnist and newsletter editor of The Bill Inmon Channel on the Business Intelligence Network. * Long overdue comprehensive coverage of the implementation of technology and tools that enable the new generation of the DW: metadata, temporal data, ETL, unstructured data, and data quality control.

This book helps you to understand Snowflake's unique architecture and ecosystem that places it at the forefront of cloud data warehouses. The recipes present in this book will enable you to develop proficiency in managing data on Snowflake and learn Snowflake's novel features such as data sharing, cloning, and time travel.

Build a modern data warehouse on Microsoft's Azure Platform that is flexible, adaptable, and fast—fast to snap together, reconfigure, and fast at delivering results to drive good decision making in your business. Gone are the days when data warehousing projects were lumbering dinosaur-style projects that took forever, drained budgets, and produced business intelligence (BI) just in time to tell you what to do 10 years ago. This book will show you how to assemble a data warehouse solution like a jigsaw puzzle by connecting specific Azure technologies that address your own needs and bring value to your business. You will see how to implement a range of architectural patterns using batches, events, and streams for both data lake technology and SQL databases. You will discover how to manage metadata and automation to accelerate the development of your warehouse while establishing resilience at every level. And you will know how to feed downstream analytic solutions such as Power BI and Azure Analysis Services to empower data-driven decision making that drives your business forward toward a pattern of success. This book teaches you how to employ the Azure platform in a strategy to dramatically improve implementation speed and flexibility of data warehousing systems. You will know how to make correct decisions in design, architecture, and infrastructure such as choosing which type of SQL engine (from at least three options) best meets the needs of your organization. You also will learn about ETL/ELT structure and the vast number of accelerators and patterns that can be used to aid implementation and ensure resilience. Data warehouse developers and architects will find this book a tremendous resource for moving their skills into the future through cloud-based implementations. What You Will Learn Choose the appropriate Azure SQL engine for implementing a given data warehouse Develop smart, reusable ETL/ELT processes that are resilient and easily maintained Automate mundane development tasks through tools such as PowerShell Ensure consistency of data by creating and enforcing data contracts Explore streaming and event-driven architectures for data ingestion Create advanced staging layers using Azure Data Lake Gen 2 to feed your data warehouse Who This Book Is For Data warehouse or ETL/ELT developers who wish to implement a data warehouse project in the Azure cloud, and developers currently working in on-premise environments who want to move to the cloud, and for developers with Azure experience looking to tighten up their implementation and consolidate their knowledge

Agile Data Warehouse Design is a step-by-step guide for capturing data warehousing/business intelligence (DW/BI) requirements and turning them into high performance dimensional models in the most direct way: by modelstorming (data modeling | brainstorming) with BI stakeholders. This book describes BEAM, an agile approach to dimensional modeling, for improving communication between data warehouse designers, BI stakeholders and the whole DW/BI development team. BEAM provides tools and techniques that will encourage DW/BI designers and developers to move away from their keyboards and entity relationship based tools and model interactively with their colleagues. The result is everyone thinks dimensionally from the outset! Developers understand how to efficiently implement dimensional modeling solutions. Business stakeholders feel ownership of the data warehouse they have created, and can already imagine how they will use it to answer their business questions. Within this book, you will learn: Agile dimensional modeling using Business Event Analysis & Modeling (BEAM) Modelstorming: data modeling that is quicker, more inclusive, more productive, and frankly more fun! Telling dimensional data stories using the 7Ws (who, what, when, where, how many, why and how) Modeling by example not abstraction; using data story themes, not crow's feet, to describe detail Storyboarding the data warehouse to discover conformed dimensions and plan iterative development Visual modeling: sketching timelines, charts and grids to model complex process measurement - simply Agile design documentation: enhancing star schemas with BEAM dimensional shorthand notation Solving difficult DW/BI performance and usability problems with proven dimensional design patterns Lawrence Corr is a data warehouse designer and educator. As Principal of DecisionOne Consulting, he helps clients to review and simplify their data warehouse designs, and advises vendors on visual data modeling techniques. He regularly teaches agile dimensional modeling courses worldwide and has taught dimensional DW/BI skills to thousands of students. Jim Stagnitto is a data warehouse and master data management architect specializing in the healthcare, financial services, and information service industries. He is the founder of the data warehousing and data mining consulting firm Llumino.

A thorough update to the industry standard for designing, developing, and deploying data warehouse and business intelligence systems The world of data warehousing has changed remarkably since the first edition of The Data Warehouse Lifecycle Toolkit was published in 1998. In that time, the data warehouse industry has reached full maturity and acceptance, hardware and software have made staggering advances, and the techniques promoted in the premiere edition of this book have been adopted by nearly all data warehouse vendors and practitioners. In addition, the term "business intelligence" emerged to reflect the mission of the data warehouse: wrangling the data out of source systems, cleaning it, and delivering it to add value to the business. Ralph Kimball and his colleagues have refined the original set of Lifecycle methods and techniques based on their consulting and training experience. The authors understand first-hand that a data warehousing/business intelligence (DW/BI) system needs to change as fast as its surrounding organization evolves. To that end, they walk you through the detailed steps of designing, developing, and deploying a DW/BI system. You'll learn to create adaptable systems that deliver data and analyses to business users so they can make better business decisions. Managing Data in Motion describes techniques that have been developed for significantly reducing the complexity of managing system interfaces and enabling scalable architectures. Author April Reeve brings over two decades of experience to present a vendor-neutral approach to moving data between computing environments and systems. Readers will learn the techniques, technologies, and best practices for managing the passage of data between computer systems and integrating disparate data together in an enterprise environment. The average enterprise's computing environment is comprised of hundreds to thousands of computer systems that have been built, purchased, and acquired over time. The data from these various systems needs to be integrated for reporting and analysis, shared for business transaction processing, and converted from one format to another when old systems are replaced and new systems are acquired. The management of the "data in motion" in organizations is rapidly becoming one of the biggest concerns for business and IT management. Data warehousing and conversion, real-time data integration, and cloud and "big data" applications are just a few of the challenges facing organizations and businesses today. Managing Data in Motion tackles these and other topics in a style easily understood by business and IT managers as well as programmers and architects. Presents a vendor-neutral overview of the different technologies and techniques for moving data between computer systems including the emerging solutions for unstructured as well as structured data types Explains, in non-technical terms, the architecture and components required to perform data integration Describes how to reduce the complexity of managing system interfaces and enable a scalable data architecture that can handle the dimensions of "Big Data"

G geared to IT professionals eager to get into the all-important field of data warehousing, this book explores all topics needed by those who design and implement data warehouses. Readers will learn about planning requirements, architecture, infrastructure, data preparation, information delivery, implementation, and maintenance. They'll also find a wealth of industry examples garnered from the author's 25 years of experience in designing and implementing databases and data warehouse applications for major corporations. Market: IT Professionals, Consultants.

The vast majority of software applications use relational databases that virtually every application developer must work with. This book introduces you to database design, whether you're a DBA or database developer. You'll discover what databases are, their goals, and why proper design is necessary to achieve those goals. Additionally, you'll master how to structure the database so it gives good performance while minimizing the chance for error. You will learn how to decide what should be in a database to meet the application's requirements.

[Concepts, Architectures, and Solutions](#)

[Database Systems: Design, Implementation, & Management](#)

[Managing Data in Motion](#)

[The Data Warehouse Lifecycle Toolkit](#)

[Practical Techniques for Extracting, Cleaning, Conforming, and Delivering Data](#)

[Data Architecture: A Primer for the Data Scientist](#)

[Second International Conference on Advances in Communication, Network, and Computing, CNC 2011, Bangalore, India, March 10-11, 2011. Proceedings](#)

[Exam 70-463 Implementing a Data Warehouse with Microsoft SQL Server 2012](#)

[The Microsoft® Data Warehouse Toolkit](#)

[Mastering Data Warehouse Design](#)

[The Modern Data Warehouse in Azure](#)

[Building with Speed and Agility on Microsoft's Cloud Platform](#)

A cutting-edge response to Ralph Kimball's challenge to the data warehouse community that answers some tough questions about the effectiveness of the relational approach to data warehousing Written by one of the best-known exponents of the Bill Inmon approach to data warehousing Addresses head-on the tough issues raised by Kimball and explains how to choose the best modeling technique for solving common data warehouse design problems Weighs the pros and cons of relational vs. dimensional modeling techniques Focuses on tough modeling problems, including creating and maintaining keys and modeling calendars, hierarchies, transactions, and data quality

This book constitutes the refereed proceedings of the Second International Conference on Advances in Communication, Network, and Computing, CNC 2011, held in Bangalore, India, in March 2011. The 41 revised full papers, presented together with 50 short papers and 39 poster papers, were carefully reviewed and selected for inclusion in the book. The papers feature current research in the field of Information Technology, Networks, Computational Engineering, Computer and Telecommunication Technology, ranging from theoretical and methodological issues to advanced applications.

Readers gain a solid foundation in database design and implementation with the practical and easy-to-understand approach in DATABASE SYSTEMS: DESIGN, IMPLEMENTATION, AND MANAGEMENT, 12E. Filled with diagrams, illustrations, and tables, this market-leading text provides in-depth coverage of database design. Readers learn the key to successful database implementation: proper design of databases to fit within a larger strategic view of the data environment. Renowned for its clear, straightforward writing style, this text provides an outstanding balance of theory and practice. Updates include the latest coverage of cloud data services and a new chapter on Big Data Analytics and NoSQL, including related Hadoop technologies. In addition, new review questions, problem sets, and cases offer multiple opportunities to test understanding and develop useful design skills. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

Here is the ideal field guide for data warehousing implementation. This book first teaches you how to build a data warehouse, including defining the architecture, understanding the methodology, gathering the requirements, designing the data models, and creating the databases. Coverage then explains how to populate the data warehouse and explores how to present data to users using reports and multidimensional databases and how to use the data in the data warehouse for business intelligence, customer relationship management, and other purposes. It also details testing and how to administer data warehouse operation.

SQL Server Integration Services Design Patterns is newly-revised for SQL Server 2014, and is a book of recipes for SQL Server Integration Services (SSIS). Design patterns in the book help to solve common problems encountered when developing data integration solutions. The patterns and solution examples in the book increase your efficiency as an SSIS developer, because you do not have to design and code from scratch with each new problem you face. The book's team of expert authors take you through numerous design patterns that you'll soon be using every day, providing the thought process and technical details needed to support their solutions. SQL Server Integration Services Design Patterns goes beyond the surface of the immediate problems to be solved, delving into why particular problems should be solved in certain ways. You'll learn more about SSIS as a result, and you'll learn by practical example. Where appropriate, the book provides examples of alternative patterns and discusses when and where they should be used. Highlights of the book include sections on ETL Instrumentation, SSIS Frameworks, Business Intelligence Markup Language, and Dependency Services. Takes you through solutions to common data integration challenges Provides examples involving Business Intelligence Markup Language Teaches SSIS using practical examples

[Beginning Database Design Solutions](#)

[Data Warehouse Design Solutions](#)

[Simple Biz 360](#)

[The Complete Guide to Dimensional Modeling](#)

[Building a Scalable Data Warehouse with Data Vault 2.0](#)

[The Definitive Guide to Dimensional Modeling](#)

[Solutions from the Experts](#)

[With SQL Server® 2005 and the Microsoft Business Intelligence Toolset](#)

[Designing a Data Warehouse](#)

[Building the Data Warehouse](#)

[Corporate Information Factory](#)

[Building and Maintaining a Data Warehouse](#)